

Chapter 3

Layer 2 VPN Configuration Guidelines

To configure Layer 2 virtual private network (VPN) functionality, you must enable Layer 2 VPN support on the provider edge (PE) router. You must also configure PE routers to distribute routing information to the other PE routers in the VPN and configure the circuits between the PE routers and the customer edge (CE) routers.

Each Layer 2 VPN is configured under a routing instance of type l2vpn. An l2vpn routing instance can transparently carry Layer 3 traffic across the service provider's network. As with other routing instances, all logical interfaces belonging to a Layer 2 VPN routing instance are listed under that instance.

The configuration of the CE routers is not relevant to the service provider. The CE routers only need to provide appropriate Layer 2 circuits (with appropriate circuit identifiers, such as Data-Link Connection Identifier [DLCI], Virtual Path Identifier/Virtual Circuit Identifier [VPI/VCI], or Virtual Local Area Network Identifier [VLAN ID]) to send traffic to the PE router.

To configure Layer 2 VPNs, you include statements at the [edit routing-instances *routing-instance-name*] hierarchy level:

```
[edit]
routing-instances {
  routing-instance-name {
    description text;
    instance-type l2vpn;
    interface interface-name;
    route-distinguisher (as-number:id | ip-address:id);
    vrf-export [ policy-name ];
    vrf-import [ policy-name ];
    protocols {
      l2vpn {
        encapsulation type
        traceoptions {
          file filename <replace> <size size> <files number> <nostamp>;
          flag flag <flag-modifier> <disable>;
        }
        site site-name {
          site-identifier identifier;
          interface interface-name {
            remote-site-id remote-site-ID;
          }
        }
      }
    }
  }
}
```

For Layer 2 VPNs, only some of the statements in the [edit routing-instances] hierarchy are valid. For the full hierarchy, see the *JUNOS Internet Software Configuration Guide: Routing and Routing Protocols*.

In addition to these statements, you must configure Multiprotocol Label Switching (MPLS) label-switched paths (LSPs) between the PE routers, Internal Border Gateway Protocol (IBGP) sessions between the PE routers, and an Interior Gateway Protocol (IGP) on the PE and provider routers.

By default, Layer 2 VPNs are disabled.

This chapter describes the following tasks for configuring Layer 2 VPNs:

Configure MPLS LSPs between the PE Routers on page 10

Configure an IGP on the PE Routers on page 13

Configure an IBGP Session between PE Routers on page 14

Configure Routing Instances for Layer 2 VPNs on the PE Routers on page 14

Configure the Connections to the Local Site on page 19

Configure MPLS LSPs between the PE Routers

For Layer 2 VPNs to function, you must configure MPLS LSPs between the PE routers. You can do one of the following:

Configure MPLS LSPs using LDP on page 10

Configure MPLS LSPs Using RSVP on page 12

Configure MPLS LSPs using LDP

To use the Label Distribution Protocol (LDP) to configure the MPLS LSPs, perform the following steps on the PE and provider routers:

1. Configure LDP on the interfaces in the core of the service provider's network by including the `ldp` statement at the [edit protocols] hierarchy level. You need to configure LDP only on the interfaces between PE routers or between PE and provider routers. You can think of these as the “core-facing” interfaces.

```
[edit]
protocols {
  ldp {
    interface interface-name;
  }
}
```

2. Configure the MPLS address family on the interfaces on which you enabled LDP (the interfaces you configured in Step 1):

```
[edit]
interfaces {
  interface-name {
    unit logical-unit-number {
      family mpls;
    }
  }
}
```

Specify the interface name in the format *type-fpc/pic/port*.

3. Configure OSPF or IS-IS on each PE and provider router. You configure these protocols at the master instance of the routing protocol, not within the routing instance used for the VPN.

To configure OSPF, include the `ospf` statement at the `[edit protocols]` hierarchy level. At a minimum, you must configure a backbone area on at least one of the router's interfaces.

```
[edit]
protocols {
  ospf {
    area 0.0.0.0 {
      interface interface-name;
    }
  }
}
```

To configure IS-IS, include the `isis` statement at the `[edit protocols]` hierarchy level and configure the loopback interface and ISO family at the `[edit interfaces]` hierarchy level. At a minimum, you must enable IS-IS on the router, configure a network entity title (NET) on one of the router's interfaces (preferably the loopback interface, `lo0`), and configure the ISO family on all interfaces on which you want IS-IS to run. When you enable IS-IS, Level 1 and Level 2 are enabled by default. The following is the minimum IS-IS configuration. In the address statement, *address* is the NET.

```
[edit]
interfaces {
  lo0 {
    unit logical-unit-number {
      family iso {
        address address;
      }
    }
  }
  type-fpc/pic/port {
    unit logical-unit-number {
      family iso;
    }
  }
}
protocols {
  isis {
    interface all;
  }
}
```

For detailed information about how to configure LDP, see the *JUNOS Internet Software Configuration Guide: MPLS Applications*. For more information about configuring OSPF and IS-IS, see the *JUNOS Internet Software Configuration Guide: Routing and Routing Protocols*.

Configure MPLS LSPs Using RSVP

To configure the MPLS LSPs using RSVP, perform the following steps:

1. On each PE router, configure traffic engineering. To do this, you must configure an IGP that supports traffic engineering (either IS-IS or OSPF) and enable traffic engineering support for that protocol.

To enable OSPF traffic engineering support, include the traffic-engineering statement at the [edit protocols ospf] hierarchy level:

```
[edit protocols ospf]
traffic-engineering;
```

For IS-IS, traffic engineering support is enabled by default.

2. On each PE and provider router, enable RSVP on the router interfaces that participate in the label-switched path (LSP). On the PE router, these are the interfaces that are the ingress and egress points to the LSP. On the provider router, these are the interfaces that connect the LSP between the PE routers.

To configure RSVP on the PE and provider routers, include the interface statement at the [edit rsvp] hierarchy level. Include one interface statement for each interface on which you are enabling RSVP.

```
[edit]
rsvp {
  interface interface-name;
  interface interface-name;
}
```

3. On each PE router, configure an MPLS LSP to the PE router that is the LSP's egress point. To do this, include the label-switched-path and interface statements at the [edit mpls] hierarchy level.

```
[edit]
mpls {
  label-switched-path lsp-path-name {
    to ip-address;
  }
  interface interface-name;
}
```

In the to statement, specify the address of the LSP's egress point, which is an address on the remote PE router.

In the interface statement, specify the name of the interface (both the physical and logical portions). Include one interface statement for the interface associated with the LSP.

When you configure the same interface at the [edit interfaces] hierarchy level, you must also configure family mpls and family inet when configuring the logical interface:

```
[edit interfaces]
interface-name {
  unit logical-unit-number {
    family inet;
    family mpls;
  }
}
```

4. On all provider routers that participate in the LSP, enable MPLS by including the interface statement at the [edit mpls] hierarchy level. Include one interface statement for each connection to the LSP.

```
[edit]
mpls {
  interface interface-name;
  interface interface-name;
}
```

5. Enable MPLS on the interface between the PE and CE routers by including the interface statement at the [edit mpls] hierarchy level. Doing this allows the PE router to assign an MPLS label to traffic entering the LSP or to remove the label from traffic exiting the LSP.

```
[edit]
mpls {
  interface interface-name;
}
```

For information about configuring RSVP or MPLS, see the *JUNOS Internet Software Configuration Guide: MPLS Applications*.

Configure an IGP on the PE Routers

To allow PE routers to exchange routing information, you must configure either an IGP or static routes on these routers. You configure the IGP on the master instance of the routing protocol process at the [edit protocols] hierarchy level, not within the routing instance used for the Layer 2 VPN—that is, not at the [edit routing-instances] hierarchy level.

When you configure the PE router, do not configure any summarization of the PE router's loopback addresses at the area boundary. Each PE router's loopback address should appear as a separate route.

For information about configuring routing protocols and static routes, see the *JUNOS Internet Software Configuration Guide: Routing and Routing Protocols*.

Configure an IBGP Session between PE Routers

You must configure an IBGP session between PE routers to allow these routers to exchange information about Layer 2 VPNs, particularly information about sites connected to Layer 2 VPNs. The PE routers rely on this information to determine which labels to use for traffic destined for remote sites. To enable an IBGP session between the PE routers, include the family l2vpn statement when configuring IBGP in the master instance:

```
[edit protocols bgp]
bgp {
  group group-name {
    type internal;
    local-address ip-address;
    family l2vpn {
      unicast;
    }
    neighbor ip-address;
  }
}
```

The family l2vpn statement indicates that the IBGP session is for the Layer 2 VPN.

The IP address in the local-address statement is the same as the address configured in the to statement at the [edit protocols mpls label-switched-path *lsp-path-name*] hierarchy level on the remote PE router. The IBGP session uses this address as the source in the peering session.

The IP address in the neighbor statement is the loopback address of the neighboring PE router. If you are using RSVP signaling, this IP address is the same address you specify in the to statement at the [edit mpls label-switched-path *lsp-path-name*] hierarchy level when you configure the MPLS LSP.

Configure Routing Instances for Layer 2 VPNs on the PE Routers

To configure routing instances for Layer 2 VPNs, include the routing-instances statement at the [edit] hierarchy level. You configure Layer 2 VPN routing instances only on the PE routers.

Configure the routing instance as follows:

```
[edit]
routing-instances {
  routing-instance-name {
    description text;
    instance-type l2vpn;
    interface interface-name;
    route-distinguisher (as-number:id | ip-address:id);
    vrf-export [ policy-name ]
    vrf-import [ policy-name ]
  }
}
```



Note

For the Layer 2 VPN to function, you must include the instance-type, interface, route-distinguisher, vrf-export, and vrf-import statements in the routing instance configuration on the PE router.

The following sections describe how to configure Layer 2 VPN routing instances:

- Configure the Description on page 15
- Configure the Instance Type on page 15
- Configure Interfaces for Layer 2 VPN Routing on page 15
- Configure CCC Encapsulation on Interfaces on page 16
- Configure TCC Encapsulation on Interfaces on page 17
- Configure Layer 2 VPN Policing on Interfaces on page 18
- Configure the Route Distinguisher on page 19
- Configure Policy for the PE Router's VRF Table on page 19

Configure the Description

To provide a textual description for the routing instance, include the description statement at the [edit routing-instances *routing-instance-name*] hierarchy level. Enclose any descriptive text that includes spaces in quotation marks (" "). Any descriptive text you include is displayed in the output of the show route instance detail command and has no effect on the operation of the routing instance.

```
[edit routing-instances routing-instance-name]
description text;
```

Configure the Instance Type

To enable Layer 2 VPN routing on a PE router, include the instance-type statement at the [edit routing-instances *routing-instance-name*] hierarchy level, specifying the instance type as l2vpn:

```
[edit routing-instances routing-instance-name]
instance-type l2vpn;
```

Configure Interfaces for Layer 2 VPN Routing

On each PE router, you must configure the interfaces over which the Layer 2 VPN traffic travels between PE and CE routers. To do this, include the interface statement at the [edit routing-instances *routing-instance-name*] hierarchy level:

```
[edit routing-instances routing-instance-name]
interface interface-name;
```

Specify both the physical and logical portions of the interface name, in the following format:

physical.logical

For example, in at-1/2/1.2, at-1/2/1 is the physical portion of the interface name and 2 is the logical portion. If you do not specify the logical portion of the interface name, 0 is set by default.

A logical interface can be associated with only one routing instance.



Note

If you enable a routing protocol on all instances by specifying interfaces all when configuring the master instance of the protocol at the [edit protocols] hierarchy level and if you configure a specific interface for Layer 2 VPN routing at the [edit routing-instances *routing-instance-name*] hierarchy level, the latter interface statement takes precedence and the interface is used exclusively for the Layer 2 VPN.

If you explicitly configure the same interface name at both the [edit protocols] and [edit routing-instances *routing-instance-name*] hierarchy levels, when you try to commit the configuration, it will fail.

Configure CCC Encapsulation on Interfaces

You need to specify a circuit cross-connect (CCC) encapsulation type for each PE-router-to-CE-router interface running a Layer 2 VPN. This encapsulation type should match the encapsulation type configured under the routing instance. See “Configure the Encapsulation Type” on page 21 for information about how to configure the encapsulation type under the routing instance.

To configure the CCC encapsulation type, include the following statements at the [edit interfaces] hierarchy level:

```
[edit]
interfaces {
  interface name {
    encapsulation ccc-encapsulation-type;
    unit unit number {
      encapsulation ccc-encapsulation-type;
    }
  }
}
```

You configure the encapsulation type at the [edit interfaces] hierarchy level differently than you do at the [edit routing-instance] hierarchy level. For example, you specify the encapsulation as frame-relay at the [edit routing-instances] hierarchy level and as frame-relay-ccc at the [edit interfaces] hierarchy level.

You can run both standard Frame Relay and CCC Frame Relay on the same device. If you specify Frame Relay encapsulation (frame-relay-ccc) for the interface, you should also configure the encapsulation at the [edit interfaces *interface name* unit *unit-number*] hierarchy level as frame-relay-ccc. Otherwise, the logical interface unit defaults to standard Frame Relay.

The following are the CCC encapsulation types:

```
atm-aal5-ccc—ATM AAL/5
atm-cell-ccc—ATM cell
cisco-hdlc-ccc—Cisco Systems-compatible HDLC
ethernet-vlan-ccc—Ethernet VLAN
frame-relay-ccc—Frame Relay
ppp-ccc—PPP
```

Configure TCC Encapsulation on Interfaces

Translation cross-connect (TCC) encapsulation types allow you to configure a different encapsulation type at the ingress and egress of a Layer 2 VPN. For example, a CE router at the ingress of a Layer 2 VPN circuit can send traffic as Frame Relay. A CE router at the egress of that circuit can receive the traffic as ATM.

The configuration for TCC encapsulation types is similar to the configuration for CCC encapsulation types. Specify a TCC encapsulation type for each PE-router-to-CE-router interface running a Layer 2 VPN. This encapsulation type should match the encapsulation type configured under the routing instance. See “Configure the Encapsulation Type” on page 21 for information about how to configure the encapsulation type under the routing instance.

To configure the TCC encapsulation type, include the following statements at the [edit interfaces] hierarchy level:

```
[edit]
interfaces {
  interface name {
    encapsulation tcc-encapsulation-type;
    unit unit number {
      encapsulation tcc-encapsulation-type;
    }
  }
}
```

You configure the encapsulation type at the [edit interfaces] hierarchy level differently than you do at the [edit routing-instance] hierarchy level. For example, you specify the encapsulation as frame-relay at the [edit routing-instances] hierarchy level and as frame-relay-tcc at the [edit interfaces] hierarchy level.

The following are the TCC encapsulation types:

```
atm-aal5-tcc—ATM AAL/5
atm-cell-tcc—ATM cell
cisco-hdlc-tcc—Cisco Systems-compatible HDLC
ethernet-tcc—Ethernet
extended-vlan-tcc—Ethernet extended VLAN
```

```
frame-relay-tcc—Frame Relay
```

```
ppp-tcc—PPP
```

Configure Layer 2 VPN Policing on Interfaces

You can use policing to control the amount of traffic flowing over the interfaces servicing a Layer 2 VPN. If policing is disabled on an interface, all the available bandwidth on a Layer 2 VPN tunnel can be used by a single circuit cross-connect (CCC) or translational cross-connect (TCC) interface.

Layer 2 VPNs support only input policing. For more information about the policer statement, see the *JUNOS Internet Software Configuration Guide: Policy Framework*.

If you configure CCC encapsulation, then include the policer statement at the [edit interfaces *interface-name* unit *unit-number* family ccc] hierarchy level to enable Layer 2 VPN policing on an interface:

```
[edit]
interfaces interface-name {
  encapsulation encapsulation-type;
  unit 0 {
    family ccc {
      policer {
        input input-name;
      }
    }
  }
}
```

If you configure TCC encapsulation, then include the policer statement at the [edit interfaces *interface-name* unit *unit-number* family tcc] hierarchy level to enable Layer 2 VPN policing on an interface:

```
[edit]
interfaces interface-name {
  encapsulation encapsulation-type;
  unit 0 {
    family tcc {
      policer {
        input input-name;
      }
    }
  }
}
```

For information about how to configure the encapsulation type, see “Configure the Encapsulation Type” on page 21.

Configure the Route Distinguisher

Each routing instance that you configure on a PE router must have a unique route distinguisher associated with it. Layer 2 VPNs need a route distinguisher to help BGP distinguish overlapping NLRIs from different VPNs. Layer 3 VPNs need a route distinguisher for the same purpose.

We recommend that you use unique route distinguishers for each routing instance that you configure. Although you can use the same route distinguisher on all PE routers in the same Layer 2 VPN, if you use a unique route distinguisher, you can determine the PE router from which a route originated.

To configure a route distinguisher on a PE router, include the `route-distinguisher` statement at the `[edit routing-instances routing-instance-name]` hierarchy level:

```
[edit routing-instances routing-instance-name]  
route-distinguisher (as-number:number | ip-address:number);
```

The route distinguisher is a 6-byte value that you can specify in one of the following formats:

as-number:number, where *as-number* is an AS number (a 2-byte value) and *number* is any 4-byte value. The AS number can be in the range 1 through 65,535. We recommend that you use an IANA-assigned, nonprivate AS number, preferably the ISP's own or the customer's own AS number.

ip-address:number, where *ip-address* is an IP address (a 4-byte value) and *number* is any 2-byte value. The IP address can be any globally unique unicast address. We recommend that you use the address that you configure in the `router-id` statement, which is a nonprivate address in your assigned prefix range.

Configure Policy for the PE Router's VRF Table

For information about configuring the VRF table, see "Configure Policy for the PE Router's VRF Table" on page 80.

Configure the Connections to the Local Site

For each local site, the PE router advertises a set of VPN labels to the other PE routers servicing the Layer 2 VPN. The VPN labels constitute a single block of contiguous labels; however, to allow for reprovisioning, more than one such block can be advertised. Each label block consists of a label base, a range (the size of the block), and a remote site ID that identifies the sequence of remote sites that connect to the local site using this label block (the remote site ID is the first site identifier in the sequence). The encapsulation type is also advertised along with the label block.

To configure the connections to the local site on the PE router, perform the following tasks:

Configure the Local Site on page 20

Configure the Encapsulation Type on page 21

Trace Layer 2 VPN Traffic and Operations on page 21

Configure the Local Site

All of the Layer 2 circuits provisioned for a local site are listed as the set of logical interfaces (using the interface statement) within the site statement.

On each PE router, you must configure each site that has a circuit to the PE router. To do this, include the site statement at the [edit routing-instances *routing-instance-name* protocols l2vpn] hierarchy level:

```
[edit routing-instances routing-instance-name protocols l2vpn]
site site-name {
  site-identifier identifier;
  interface interface-name {
    remote-site-id remote-site-ID;
  }
}
```

You must configure the following for each site:

site—Name of the site.

site-identifier—Unsigned 16-bit number greater than zero that uniquely identifies the site.

interface—A name for the interface and, optionally, a remote site ID for remote site connections.

Under the interface statement, you can set a remote site ID, which identifies the remote site to which this interface connects. Be aware of the order of the interfaces because this determines which remote site each interface connects to. The order of the interfaces is based on the interface's site identifier.

The remote-site-id statement is optional; if omitted, the remote site ID for an interface is automatically set to 1 higher than the remote site ID for the previous interface. For example, if the first interface in the list does not have a remote site ID, its offset is set to 1. The second interface in the list has its offset set to 2, and the third has its offset set to 3. The offsets of any interfaces that follow are incremented in the same manner if you do not explicitly configure them.

The remote site ID allows for a sparse Layer 2 VPN topology. When you configure remote site IDs, each site does not have to connect to all other sites in the Layer 2 VPN, making it unnecessary to allocate circuits for all the remote sites. Remote site IDs are particularly important if you configure a topology more complicated than full mesh.

Configure the Encapsulation Type

The encapsulation type you configure at each Layer 2 VPN site varies depending on which Layer 2 protocol you choose to configure. You need to use the same protocol at each Layer 2 VPN site if you configure ethernet-vlan as the encapsulation type. You do *not* need to use the same protocol at each Layer 2 VPN site if you configure any of the following encapsulation types:

atm-aal5—ATM AAL/5

atm-cell—ATM cell

cisco-hdlc—Cisco Systems-compatible HDLC

frame-relay—Frame Relay

ppp—PPP

Note that if you configure different protocols at your Layer 2 VPN sites, you need to configure a TCC encapsulation type. For more information, see “Configure TCC Encapsulation on Interfaces” on page 17.

To configure the Layer 2 protocol accepted by the PE router, specify the encapsulation type by including the encapsulation statement at the [edit routing-instances *routing-instance-name* protocols l2vpn] hierarchy level:

```
[edit routing-instances routing-instance-name protocols l2vpn]
encapsulation type
```

Trace Layer 2 VPN Traffic and Operations

To trace Layer 2 VPN protocol traffic, you can specify options in the Layer 2 VPN traceoptions statement at the [edit routing-instances *routing-instance-name* protocols l2vpn] hierarchy level:

```
[edit routing-instances routing-instance-name protocols l2vpn]
traceoptions {
  file filename <replace> <size size> <files number> <nostamp>;
  flag flag <flag-modifier> <disable>;
}
```

The following trace flags display the operations associated with Layer 2 VPNs. Each can carry one or more of the following modifiers:

all—All Layer 2 VPN tracing options

connections—Layer 2 VPN connections (events and state changes)

error—Error conditions

nlri—Layer 2 VPN advertisements received or sent using BGP

route—Trace routing information

topology—Layer 2 VPN topology changes due to reconfiguration or due to advertisements received from other PE routers using BGP

Disable Normal TTL Decrementing for VPNs

To diagnose networking problems related to VPNs (Layer 2 or Layer 3), it can be useful to disable normal TTL decrementing. In JUNOS, you can do this with the `no-propagate-ttl` and `no-decrement-ttl` statements. However, when tracing VPN traffic, only the `no-propagate-ttl` statement is effective.

For the `no-propagate-ttl` statement to have an effect on VPN behavior, you need to clear the PE router-to-PE router BGP session, or disable and then enable the VPN routing instance.

For more information about the `no-propagate-ttl` and `no-decrement-ttl` statements, see the *JUNOS Internet Software Configuration Guide: MPLS Applications*.